



Microsoft

Exam Questions DP-900

Microsoft Azure Data Fundamentals

NEW QUESTION 1

Which Azure SQL offering supports automatic database scaling and automatic pausing of the database during inactive periods?

- A. Azure SQL Database Hyperscale
- B. Azure SQL managed instance
- C. Azure SQL Database serveless
- D. Azure SQL Database elastic pod

Answer: B

NEW QUESTION 2

A bar chart showing year-to date sales by region is an example of which type of analytics?

- A. descriptive
- B. diagnostic
- C. predictive
- D. prescriptive

Answer: B

NEW QUESTION 3

To complete the sentence, select the appropriate option in the answer area.

A visualization that illustrates a university’s current student enrollment per department is an example of

▼

cognitive

descriptive

predictive

prescriptive

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text, table Description automatically generated

Generally speaking, data analytics comes in four types:

- > Descriptive, to answer the question: What’s happening?
- > Diagnostic, to answer the question: Why’s happening?
- > Predictive, to answer the question: What will happen?
- > Prescriptive, to answer the question: What actions should we take? Company name Description automatically generated



Descriptive

What's happening?



Diagnostic

Why's happening?



Predictive

What will happen?



Prescriptive

What actions should we take?

NEW QUESTION 4

Select the answer that correctly completes the sentence.

Answer Area

Column family databases

▼

Column family databases

Document databases

Graph databases

Key-value stores

_____ natively support the analysis of relationships between entities.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area



natively support the analysis of relationships between entities.

NEW QUESTION 5

You have an Azure Cosmos DB account that uses the Core (SQL) API.

Which two settings can you configure at the container level? Each correct answer presents a complete solution. (Choose two.)

NOTE: Each correct selection is worth one point.

- A. the throughput
- B. the read region
- C. the partition key
- D. the API

Answer: AC

NEW QUESTION 6

Your company is designing a data store that will contain student data. The data has the following format.

StudentNumber	StudentInformation
7634634	First name: Ben Last: Smith Preferred Name: Benjamin
7634634	First Name: Dominik Last Name: Paiha Email Address: dpaiha@contoso.com MCP ID: 931817
7634636	First Name: Reshma Last Name: Patel Phone number: 514-555-1101
7634637	First Name: Yun-Feng Last Name: Peng

Which type of data store should you use?

- A. graph
- B. key/value
- C. object
- D. columnar

Answer: D

NEW QUESTION 7

What are two uses of data visualization? Each correct answer present a complete solution. NOTE: Each correct selection is worth one point.

- A. Represent trends and patterns over time.
- B. Communicate the significance of data.
- C. implement machine learning to predict future values.
- D. Consistently implement business logic across reports

Answer: AB

Explanation:

Data visualization is a key component in being able to gain insight into your data. It helps make big and small data easier for humans to understand. It also makes it easier to detect patterns, trends, and outliers in groups of data.

Data visualization brings data to help you find key business insights quickly and effectively.

NEW QUESTION 8

Which scenario is an example of a streaming workload?

- A. sending transactions that are older than a month to an archive
- B. sending transactions daily from point of sale (POS) devices
- C. sending telemetry data from edge devices
- D. sending cloud infrastructure metadata every 30 minutes

Answer: C

NEW QUESTION 9

Which type of database contains nodes and edges?

- A. graph
- B. columnar
- C. key/value
- D. time series

Answer: A

NEW QUESTION 10

To complete the sentence, select the appropriate option in the answer area.

Relational data is stored in

a file system as unstructured data.

a hierarchal folder structure.

a tabular form of rows and columns.

comma-separated value (CSV) files.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

NEW QUESTION 10

To complete the sentence, select the appropriate option in the answer area.

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.

queries and their results in order to protect data in transit.

the database to protect data at rest.

the server to protect data at rest.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.

queries and their results in order to protect data in transit.

the database to protect data at rest.

the server to protect data at rest.

NEW QUESTION 14

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

An elastic pool enables multiple Azure SQL databases to share resources such as memory, data storage space, and processing power.

NEW QUESTION 16

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Batch processing can output data to a file store	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a relational database	<input type="radio"/>	<input type="radio"/>
Batch processing can output data to a NoSQL database	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
Big data solutions often use long-running batch jobs to filter, aggregate, and otherwise prepare the data for analysis. Usually these jobs involve reading source files from scalable storage (like HDFS, Azure Data Lake Store, and Azure Storage), processing them, and writing the output to new files in scalable storage.
Box 2: No
Box 3: No

NEW QUESTION 21

Which type of Azure resource supports the serverless configuration of an Azure SQL database?

- A. Azure SQL Managed Instance
- B. SQL Server on Azure Virtual Machines
- C. a single database in Azure SQL Database
- D. an Azure SQL Database elastic pool

Answer: C

NEW QUESTION 25

What is a characteristic of a non-relational database?

- A. a fixed schema
- B. self-describing entities
- C. full support for Transact-SQL

Answer: C

NEW QUESTION 28

Your company is designing an application that will write a high volume of JSON data and will have an application-defined schema.
Which type of data store should you use?

- A. key/value
- B. document
- C. graph
- D. columnar

Answer: A

Explanation:

A key/value store associates each data value with a unique key.
An application can store arbitrary data as a set of values. Any schema information must be provided by the application. The key/value store simply retrieves or stores the value by key.

NEW QUESTION 32

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<input type="radio"/>	<input type="radio"/>
Partition keys are used in Azure Cosmos DB to optimize queries.	<input type="radio"/>	<input type="radio"/>
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: No
The API determines the type of account to create. Azure Cosmos DB provides five APIs: Core (SQL) and MongoDB for document data, Gremlin for graph data, Azure Table, and Cassandra. Currently, you must create a separate account for each API.

Box 2: Yes
Azure Cosmos DB uses partitioning to scale individual containers in a database to meet the performance needs of your application. In partitioning, the items in a container are divided into distinct subsets called logical partitions. Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

Box 3: No
Logical partitions are formed based on the value of a partition key that is associated with each item in a container.

NEW QUESTION 33

Your company needs to ensure that in Azure virtual machine can connect to Azure SQL databases without exposing the databases to the internet. What should you use?

- A. Azure DNS
- B. Azure Application Gateway
- C. Azure Private link
- D. Azure Traffic Manager

Answer: C

Explanation:

Azure Private Link enables you to access Azure PaaS Services (for example, Azure Storage and SQL Database) and Azure hosted customer-owned/partner services over a private endpoint in your virtual network.
Traffic between your virtual network and the service travels the Microsoft backbone network. Exposing your service to the public internet is no longer necessary

NEW QUESTION 38

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A non-relational database

stores data without a fixed schema

NEW QUESTION 41

Match the Azure Cosmos DB APIs to the appropriate data structures.
To answer, drag the appropriate API from the column on the left to its data structure on the right. Each API may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Graph data

JSON documents

Key/value data

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

APIs

Cassandra API

Gremlin API

MongoDB API

Table API

Answer Area

Gremlin API

Graph data

MongoDB API

JSON documents

Table API

Key/value data

NEW QUESTION 44

To complete the sentence, select the appropriate option in the answer area.

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

enable the hierarchical namespace.

set Account kind to BlobStorage.

set Performance to Premium.

set Replication to Read-access geo-redundant storage (RA-GRS).

- A. Mastered
 B. Not Mastered

Answer: A

Explanation:

To configure an Azure Storage account to support both security at the folder level and atomic directory manipulation,

enable the hierarchical namespace.

set Account kind to BlobStorage.

set Performance to Premium.

set Replication to Read-access geo-redundant storage (RA-GRS).

NEW QUESTION 47

What are three characteristics of an Online Transaction Processing (OLTP) workload? Each correct answer presents a complete solution NOTE: Each correct selection is worth one point.

- A. denormalized data
 B. schema defined in a database
 C. schema defined when reading unstructured data from a database
 D. normalized data
 E. light writes and heavy reads

F. heavy writes and moderate reads

Answer: ABF

NEW QUESTION 52

Match the tools to the appropriate descriptions.

To answer, drag the appropriate tool from the column on the left to its description on the right. Each tool may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Tools

Azure Data Studio

Microsoft SQL Server Data Tools (SSDT)

Microsoft SQL Server Management Studio (SSMS)

Microsoft Visual Studio Code

Answer Area

ToolA graphical tool for managing Azure SQL databases and viewing SQL execution plans

ToolA Microsoft SQL Server extension that supports connections to SQL Server and provides a rich editing experience for T-SQL

ToolA lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files

ToolA development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Microsoft SQL Server Management Studio (SSMS)

A graphical tool for managing SQL Server or Azure SQL databases that supports access, configuration, management, and administration tasks.

Microsoft Visual Studio Code

A lightweight source code editor with an mssql extension that supports connections to SQL Server and a rich editing experience for T-SQL.

Azure Data Studio

A lightweight editor that can run on-demand SQL queries and view and save results as text, JSON, or Microsoft Excel files.

Microsoft SQL Server Data Tools (SSDT)

A development tool for building Azure SQL databases, Microsoft SQL Server relational databases, SQL Server Analysis Services (SSAS) data models, SQL Server Integration Services (SSIS) packages, and SQL Server Reporting Services (SSRS) reports.

NEW QUESTION 54

You have a transactional application that stem data in an Azure SQL managed instance. When should you implement a read-only database replica?

- A. You need to generate reports without affecting the transactional workload.
- B. You need to audit the transactional application.
- C. You need to implement high availability in the event of a regional outage
- D. You need to improve the recovery point objective (RPO).

Answer: A

Explanation:

Use read-only replicas to offload read-only query workloads.

NEW QUESTION 55

Which storage solution supports role-based access control (RBAC) at the file and folder level?

- A. Azure Disk Storage
- B. Azure Data Lake Storage
- C. Azure Blob storage
- D. Azure Queue storage

Answer: B

NEW QUESTION 56

Your company is designing a database that will contain session data for a website. The data will include notifications, personalization attributes, and products that are added to a shopping cart.

Which type of data store will provide the lowest latency to retrieve the data?

- A. key/value

- B. graph
- C. columnar
- D. document

Answer: C

NEW QUESTION 57

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input type="radio"/>	<input type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input type="radio"/>
Azure Data Studio can be used to restore a database.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure Data Studio can be used to query an Azure SQL database from a device that runs macOS.	<input type="radio"/>	<input checked="" type="radio"/>
Microsoft SQL Server Management Studio (SSMS) enables users to create and use SQL notebooks.	<input type="radio"/>	<input checked="" type="radio"/>
Azure Data Studio can be used to restore a database.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 60

What Is a benefit of hosting a database on Azure SQL managed instance as compared to an Azure SQL database?

- A. native support for cross-database queries and transactions
- B. built-in high availability
- C. system-Initiated automatic backups
- D. support for encryption at rest

Answer: A

NEW QUESTION 63

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

In a data warehousing workload, data _____ from multiple sources is combined in a single location.

NEW QUESTION 65

Which two Azure services can be used to provision Apache Spark clusters? Each correct answer presents a complete solution. (Choose two.)
NOTE: Each correct selection is worth one point.

- A. Azure Time Series Insights
- B. Azure HDInsight
- C. Azure Databricks
- D. Azure Log Analytics

Answer: BC

NEW QUESTION 67

To complete the sentence, select the appropriate option in the answer area.

To configure an Azure Storage account to support access control lists that have object-level permissions,

enable the hierarchical namespace.

set Account kind to BlobStorage.

set Performance to Premium.

set Replication to Read-access geo-redundant storage (RA-GRS).

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Table Description automatically generated

A key mechanism that allows Azure Data Lake Storage Gen2 to provide file system performance at object storage scale and prices is the addition of a hierarchical namespace. This allows the collection of objects/files within an account to be organized into a hierarchy of directories and nested subdirectories in the same way that the file system on your computer is organized. With a hierarchical namespace enabled, a storage account becomes capable of providing the scalability and cost-effectiveness of object storage, with file system semantics that are familiar to analytics engines and frameworks.

NEW QUESTION 71

You need to gather real-time telemetry data from a mobile application. Which type of workload describes this scenario?

- A. Online Transaction Processing (OLTP)
- B. batch
- C. massively parallel processing (MPP)
- D. streaming

Answer: D

NEW QUESTION 72

Which three objects can be added to a Microsoft Power BI dashboard? Each correct answer presents a complete solution. (Choose three.)
NOTE: Each correct selection is worth one point.

- A. a report page
- B. a Microsoft PowerPoint slide
- C. a visualization from a report
- D. a dataflow
- E. a text box

Answer: ACE

NEW QUESTION 75

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can copy a dashboard between Microsoft Power BI workspaces.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BI dashboard can only display visualizations from a single dataset.	<input type="radio"/>	<input type="radio"/>
A Microsoft Power BY dashboard can display visualizations from a Microsoft Excel workbook.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application Description automatically generated

Box 1: No

You can duplicate a dashboard. The duplicate ends up in the same Power BI workspace.
There is no current functionality that allows you to move reports from one workspace to another.

Box 2: No

Box 3: Yes

NEW QUESTION 77

To complete the sentence, select the appropriate option in the answer area.

An extract, transform, and load (ETL) process requires

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store.

that the data target be a relational database.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

An extract, transform, and load (ETL) process requires

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store.

that the data target be a relational database.

NEW QUESTION 78


Which database transaction property ensures that transactional changes to a database are preserved during unexpected operating system restarts?

- A. consistency
- B. durability
- C. isolation
- D. atomicity

Answer: B

NEW QUESTION 81

Your company recently reported sales from the third quarter. You have the chart shown in the following exhibit.



Which type of analysis is shown for the fourth quarter?

- A. predictive
- B. prescription
- C. descriptive
- D. diagnostic

Answer: A

NEW QUESTION 84

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A key/value data store is optimized for simple lookups.

NEW QUESTION 88

Select the answer that correctly completes the sentence.

Answer Area

A relational database must be used when strong consistency guarantees are required.

a dynamic schema is required.

data will be stored as key/value pairs.

storing large images and videos.

strong consistency guarantees are required.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A relational database must be used when strong consistency guarantees are required.

a dynamic schema is required.

data will be stored as key/value pairs.

storing large images and videos.

strong consistency guarantees are required.

NEW QUESTION 92

To complete the sentence, select the appropriate option in the answer area.

Batch workloads

	▼
process data in memory, row-by-row.	
collect and process data at most once a day.	
process data as new data is received in near real-time.	
collect data and then process the data when a condition is met.	

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Batch workloads

process data in memory, row-by-row.

collect and process data at most once a day.

process data as new data is received in near real-time.

collect data and then process the data when a condition is met.

NEW QUESTION 96

Select the answer that correctly completes the sentence.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

A relational database must be used when strong consistency guarantees are required.

NEW QUESTION 101

Match the types of data to the appropriate Azure data services.
To answer, drag the appropriate data types from the column on the left to its service on the right. Each data type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Data Types

Image files

Key/value pairs

Relationships between employees

Answer Area

Data Type

Azure Blob storage

Data Type

Azure Cosmos DB Gremlin API

Data Type

Azure Table storage

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Image files
Azure Blob storage is suitable for image files. Box 2:Key/value pairs
Azure CosmosDB table API is a key-value storage hosted in the cloud. Box 3: Relationship between employees
One-to-many relationships between business domain objects occur frequently: for example, one department has many employees. There are several ways to implement one-to-many relationships in the Azure Table service.

NEW QUESTION 103

To complete the sentence, select the appropriate option in the answer area.

An extract, load, and transform (ELT) process requires

a separate transformation engine.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store.

a data pipeline that includes a transformation engine.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Text Description automatically generated with medium confidence
With ELT, the data store used to perform the transformation is the same data store where the data is ultimately consumed.

NEW QUESTION 106

You have an e-commerce application that reads and writes data to an Azure SQL database. Which type of processing does the application use?

- A. stream processing
- B. batch processing
- C. Online Analytical Processing (OLTP)
- D. Online Transaction Processing (OLTP)

Answer: D

NEW QUESTION 108

For each of the following statements. select Yes if the statement is true. Otherwise, select No. NOTE Each correct selection is worth one point.

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<input type="radio"/>	<input type="radio"/>
Partition keys are used in Azure Cosmos DB to optimize queries.	<input type="radio"/>	<input type="radio"/>
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
The Azure Cosmos DB API is configured separately for each database in an Azure Cosmos DB account.	<input checked="" type="radio"/>	<input type="radio"/>
Partition keys are used in Azure Cosmos DB to optimize queries.	<input type="radio"/>	<input checked="" type="radio"/>
Items contained in the same Azure Cosmos DB logical partition can have different partition keys.	<input type="radio"/>	<input checked="" type="radio"/>

NEW QUESTION 111

When provisioning an Azure Cosmos DB account, which feature provides redundancy within an Azure region?

- A. multi-master replication
- B. Availability Zones
- C. automatic failover
- D. the strong consistency level

Answer: B

Explanation:

With Availability Zone (AZ) support, Azure Cosmos DB will ensure replicas are placed across multiple zones within a given region to provide high availability and resiliency to zonal failures.
Note: Azure Cosmos DB provides high availability in two primary ways. First, Azure Cosmos DB replicates data across regions configured within a Cosmos account. Second, Azure Cosmos DB maintains 4 replicas of data within a region.

NEW QUESTION 114

Select the answer that correctly completes the sentence. See the answer below:

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

An Azure Storage account must have Hierarchical namespace enabled to support Azure Data Lake Storage.



NEW QUESTION 118

Select the answer that correctly completes the sentence.

Answer Area

Microsoft SQL Server database tables store



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Microsoft SQL Server database tables store relational data.



NEW QUESTION 123

Match the security components to the appropriate scenarios.
To answer, drag the appropriate component from the column on the left to its scenario on the right. Each component may be used once, more than once, or not at all. NOTE: Each correct match is worth one point.

Components

Authentication

Firewall

Encryption

Answer Area

Prevent access to an Azure SQL database from another network.

Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.

Ensure that sensitive data never appears as plain text in an Azure SQL database.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Components

- Authentication
- Firewall
- Encryption

Answer Area

- Firewall Prevent access to an Azure SQL database from another network.
- Authentication Support Azure Active Directory (Azure AD) sign-ins to an Azure SQL database.
- Encryption Ensure that sensitive data never appears as plain text in an Azure SQL database.

NEW QUESTION 125

You have an inventory management database that contains the following table.

ProductName	Quantity
Product1	100
Product2	129
Product3	176

Which statement should you use in a SQL query to change the inventory quantity of Product1 to 270?

- A. INSERT
- B. MERGE
- C. UPDATE
- D. CREATE

Answer: C

NEW QUESTION 130

Your company needs to implement a relational database in Azure. The solution must minimize ongoing maintenance. Which Azure service should you use?

- A. Azure HDInsight
- B. Azure SQL Database
- C. Azure Cosmos DB
- D. SQL Server on Azure virtual machines

Answer: B

NEW QUESTION 134

Match the types of activities to the appropriate Azure Data Factory activities.

To answer, drag the appropriate activity type from the column on the left to its Data Factory activity on the right. Each activity type may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.

Activity Types

- Control
- Data movement
- Data transformation

Answer Area

- Copy
- Mapping data flow
- Until

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Activity Types

- Control
- Data movement
- Data transformation

Answer Area

- Data movement Copy
- Data transformation Mapping data flow
- Control Until

NEW QUESTION 135

Which component of Azure Data Factory enables you to define a sequence of activities including for each iterators?

- A. control flow
- B. linked services
- C. datasets
- D. the integration runtime

Answer: A

NEW QUESTION 138

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area		
Statements	Yes	No
Relational database tables contain columns and rows.	<input type="radio"/>	<input type="radio"/>
Indexes in a relational database describe the data types in a table.	<input type="radio"/>	<input type="radio"/>
A database view is a virtual table whose content is defined by a query.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes

Tables are database objects that contain all the data in a database. In tables, data is logically organized in a row-and-column format similar to a spreadsheet. Each row represents a unique record, and each column represents a field in the record.

Box 2: No

An index is an on-disk structure associated with a table or view that speeds retrieval of rows from the table or view.

Box 3: Yes

A view is a virtual table whose contents are defined by a query. Like a table, a view consists of a set of named columns and rows of data.

NEW QUESTION 140

To complete the sentence, select the appropriate option in the answer area.

<div><div></div><div>▼</div></div> <div><div>A heap</div><div>A stored procedure</div><div>A view</div><div>An index</div></div>	is a virtual table that contains content defined by a query.
--	--

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

<div><div></div><div>▼</div></div> <div><div>A heap</div><div>A stored procedure</div><div>A view</div><div>An index</div></div>	is a virtual table that contains content defined by a query.
--	--

NEW QUESTION 141

You are writing a set of SQL queries that administrators will use to troubleshoot an Azure SQL database. You need to embed documents and query results into a SQL notebook.

What should you use?

- A. Microsoft SQL Server Management Studio (SSMS)

- B. Azure Data Studio
- C. Azure CLI
- D. Azure PowerShell

Answer: B

NEW QUESTION 144

Match the types of workloads to the appropriate scenarios.
To answer, drag the appropriate workload type from the column on the left to its scenario on the right. Each workload type may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Workload Types

Batch

Streaming

Answer Area

Workload Type

Analyzing web traffic from the past year

Workload Type

Classifying images that were uploaded last month

Workload Type

Tracking how many people are currently using a website

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Batch
Batch processing refers to the processing of blocks of data that have already been stored over a period of time. Box 2: Streaming
Stream processing is a big data technology that allows us to process data in real-time as they arrive and detect conditions within a small period of time from the point of receiving the data. It allows us to feed data into analytics tools as soon as they get generated and get instant analytics results. Box 3: Batch

NEW QUESTION 147

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure Synapse Analytics scales storage and compute independently	<input type="radio"/>	<input type="radio"/>
Azure Synapse Analytics can be paused to reduce compute costs	<input type="radio"/>	<input type="radio"/>
An Azure Synapse Analytics data warehouse has a fixed storage capacity	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
Compute is separate from storage, which enables you to scale compute independently of the data in your system.
Box 2: Yes
You can use the Azure portal to pause and resume the dedicated SQL pool compute resources.
Pausing the data warehouse pauses compute. If your data warehouse was paused for the entire hour, you will not be charged compute during that hour.
Box 3: No
Storage is sold in 1 TB allocations. If you grow beyond 1 TB of storage, your storage account will automatically grow to 2 TBs.

NEW QUESTION 148

Match the Azure services to the appropriate requirements.
To answer, drag the appropriate service from the column on the left to its requirement on the right. Each service may be used once, more than once, or not at all.
NOTE: Each correct match is worth one point.

Services	Answer Area
Azure Data Factory	Service
Azure Data Lake Storage	Service
Azure SQL Database	Service
Azure Synapse Analytics	Service

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Azure Data Factory
 Box 2: Azure Data Lake Storage
 Azure Data Lake Storage (ADLA) now natively supports Parquet files. ADLA adds a public preview of the native extractor and outputter for the popular Parquet file format
 Box 3: Azure Synapse Analytics
 Use Azure Synapse Analytics Workspaces.

NEW QUESTION 152

You have a SQL query that combines customer data and order data. The query includes calculated columns. You need to persist the SQL query so that other users can use the query. What should you create?

- A. a table
- B. an index
- C. scalar function
- D. a view

Answer: D

Explanation:

A view is a virtual table whose contents are defined by a query. A view acts as a filter on the underlying tables referenced in the view. The query that defines the view can be from one or more tables or from other views in the current or other databases.

NEW QUESTION 157

For each of the following statements, select Yes if the statement is true. Otherwise, select No.
 NOTE: Each correct selection is worth one point.

Statements	Yes	No
Azure Databricks can consume data from Azure SQL Database	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Event Hubs	<input type="radio"/>	<input type="radio"/>
Azure Databricks can consume data from Azure Cosmos DB	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box 1: Yes
 Azure Databricks can consume data from SQL Databases using JDBC and from SQL Databases using the Apache Spark connector.
 The Apache Spark connector for Azure SQL Database and SQL Server enables these databases to act as input data sources and output data sinks for Apache Spark jobs.
 Box 2: Yes
 You can stream data into Azure Databricks using Event Hubs. Box 3: Yes
 You can run Spark jobs with data stored in Azure Cosmos DB using the Cosmos DB Spark connector. Cosmos can be used for batch and stream processing, and as a serving layer for low latency access.
 You can use the connector with Azure Databricks or Azure HDInsight, which provide managed Spark clusters on Azure.

NEW QUESTION 162

What is used to define a query in a stream processing jobs in Azure Stream Analytics?

- A. XML
- B. KQL
- C. YAML
- D. SQL

Answer: D

NEW QUESTION 163

To complete the sentence select the appropriate option in the answer area.

Answer Area

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.

queries and their results in order to protect data in transit.

the database to protect data at rest.

the server to protect data at rest.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Transparent Data Encryption (TDE) encrypts

a column to protect data at rest and in transit.

queries and their results in order to protect data in transit.

the database to protect data at rest.

the server to protect data at rest.

NEW QUESTION 168

Which two Azure services can be used to provision Spark clusters? Each correct answer presents a complete solution.

NOTE: Each correct selection is worth one point.

- A. Azure Databricks
- B. Azure Log Analytics
- C. Azure Time Series Insights
- D. Azure HDinsight

Answer: AD

NEW QUESTION 172

What is a benefit of the Azure Cosmos DB Table API as compared to Azure Table storage?

- A. supports partitioning
- B. provides resiliency if art Azure region fads
- C. provides a higher storage capacity
- D. supports a multi-master model

Answer: D

Explanation:

Multi-master support for Azure Cosmos DB is now available in all public regions.

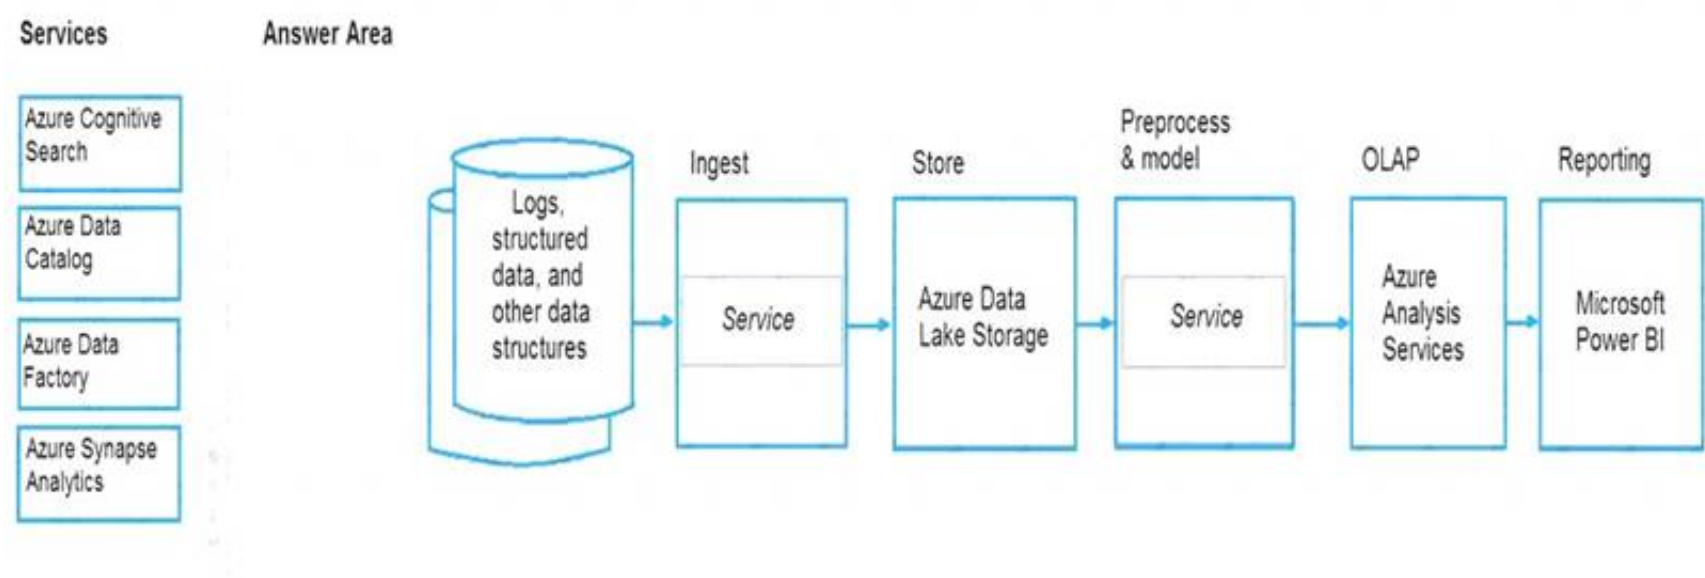
Azure CosmosDB table API is a key-value storage hosted in the cloud. It's a part of Azure Cosmos DB, that is Microsoft's multi-model database.

NEW QUESTION 174

Match the Azure services to the appropriate locations in the architecture.

To answer, drag the appropriate service from the column on the left to its location on the right. Each service may be used once, more than once, or not at all.

NOTE: Each correct match is worth one point.



- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Box Ingest: Azure Data Factory
You can build a data ingestion pipeline with Azure Data Factory (ADF). Box Preprocess & model: Azure Synapse Analytics
Use Azure Synapse Analytics to preprocess data and deploy machine learning models.

NEW QUESTION 179

Select the answer that correctly completes the sentence.

Answer Area

Online analytical processing (OLAP) systems are particularly suitable for analytical workloads because data is

optimized for transactions.

limited in size.

optimized for transactions.

pre-aggregated.

read-only.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Online analytical processing (OLAP) systems are particularly suitable for analytical workloads because data is

optimized for transactions.

limited in size.

optimized for transactions.

pre-aggregated.

read-only.

NEW QUESTION 180

You have an application that runs on Windows and requires access to a mapped drive. Which Azure service should you use?

- A. Azure Cosmos DB
- B. Azure Table storage
- C. Azure Files
- D. Azure Blob Storage

Answer: A

NEW QUESTION 184

When can you use an Azure Resource Manager template?

- A. to automate the creation of an interdependent group of Azure resources in a repeatable way
- B. to automate the creation of an interdependent and develop can deploy from the Azure portal
- C. To provision Azure subscriptions and apply policies for multi-tenant deployments
- D. To set complex permissions automatically on existing objects in an Azure subscription

Explanation:

NEW QUESTION 185

Answer Area

data that is fully processed before being loaded to the target data store.

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store

that the data target be a relational database.

- Answer: A**

Explanation:

Answer Area

data that is fully processed before being loaded to the target data store.

a matching schema in the data source and the data target.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store Content Display

that the data target be a relational database.

You have a table named Sales that contains the following data.

SalesDate	SalesAmount	ProductID
4-Apr-20	\$2,000	1
5-Apr-20	\$40	2
5-Apr-20	\$2,300	1
6-Apr-20	\$40	3
6-Apr-20	\$200	4

SalesDate	AVG(SalesAmount)
4-Apr-20	\$2,000
5-Apr-20	\$1,170
6-Apr-20	\$120

NOTE: Each correct selection is worth one point.

Values

CREATE

GROUP BY

ORDER BY

SELECT

Answer Area

Value

SalesDate, AVG(SalesAmount)

FROM Sales

Value

SalesDate

- Answer: A**

Explanation:

```
SELECT TerritoryID, AVG(Bonus)as 'Average bonus', SUM(SalesYTD) as 'YTD sales' FROM Sales.SalesPerson
GROUP BY TerritoryID;
```

NEW QUESTION 189

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Answer Area

Statements	Yes	No
Azure SQL managed instance supports cross-database queries.	<input type="radio"/>	<input type="radio"/>
Azure SQL managed instance supports user-created backups.	<input type="radio"/>	<input type="radio"/>
Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

Statements	Yes	No
Azure SQL managed instance supports cross-database queries.	<input checked="" type="radio"/>	<input type="radio"/>
Azure SQL managed instance supports user-created backups.	<input type="radio"/>	<input checked="" type="radio"/>
Azure SQL managed instance databases can be restored to SQL Server on an Azure virtual machine.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 192

For each of the following statements, select Yes if the statement is true. Otherwise, select No. NOTE: Each correct selection is worth one point.

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	<input type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input type="radio"/>	<input type="radio"/>

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Statements	Yes	No
You can use Azure Data Studio to query a Microsoft SQL Server big data cluster.	<input checked="" type="radio"/>	<input type="radio"/>
You can use Microsoft SQL Server Management Studio (SSMS) to query an Azure Synapse Analytics data warehouse.	<input checked="" type="radio"/>	<input type="radio"/>
You can use MySQL Workbench to query Azure Database for MariaDB databases.	<input checked="" type="radio"/>	<input type="radio"/>

NEW QUESTION 196

Select the answer that correctly completes the sentence.

Answer Area

In a relational database, each row in a table has

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer as below.

Answer Area

In a relational database, each row in a table has

NEW QUESTION 200

Your company has a reporting solution that has paginated reports. The reports query a dimensional model in a data warehouse. Which type of processing does the reporting solution use?

- A. stream processing
- B. Online Transaction Processing (OLTP)
- C. batch processing
- D. Online Analytical Processing (OLAP)

Answer: D

NEW QUESTION 204

Select the answer that correctly completes the sentence.

A

line chart

matrix

scatter

treemap

 is a chart of colored rectangles. The size of the rectangles represent

the relative value of each item. The chart can be hierarchical, displaying data as a set of rectangles nested within the main rectangle.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Graphical user interface, text, application, email Description automatically generated

A treemap chart divides the chart area into rectangles that represent the different levels and relative sizes of the data hierarchy.

NEW QUESTION 205

To complete the sentence, select the appropriate option in the answer area.

Answer Area

An extract, load, and transform (ELT) process requires

a data pipeline that includes a transformation engine.

a separate transformation engine.

a target data store powerful enough to transform data.

data that is fully processed before being loaded to the target data store.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Answer Area

An extract, load, and transform (ELT) process requires

- a data pipeline that includes a transformation engine.
- a separate transformation engine.
- a target data store powerful enough to transform data.
- data that is fully processed before being loaded to the target data store. 1

NEW QUESTION 210

To complete the sentence, select the appropriate option in the answer area.

Answer Area

A relational database is appropriate for scenarios that involve a high volume of

- changes to relationships between entities.
- geographically distributed writes.
- transactional writes.
- writes that have varying data structures.

- A. Mastered
- B. Not Mastered

Answer: A

Explanation:

Disadvantages of non-relational databases include: Data Consistency — non-relational databases do not perform ACID transactions.

Note: Relational databases are optimized for writes. They are optimized for consistency and availability. Advantages of relational databases include simplicity, ease of data retrieval, data integrity, and flexibility.

NEW QUESTION 215

Your company is designing a data store for internet-connected temperature sensors. The collected data will be used to analyze temperature trends. Which type of data store should you use?

- A. relational
- B. columnar
- C. graph
- D. time series

Answer: D

Explanation:

Time series data is a set of values organized by time. Time series databases typically collect large amounts of data in real time from a large number of sources.

Updates are rare, and deletes are often done as bulk operations. Although the records written to a time-series database are generally small, there are often a large number of records, and total data size can grow rapidly.

NEW QUESTION 217

You need to use Transact-SQL to query files in Azure Data Lake Storage from an Azure Synapse Analytics data warehouse. What should you use to query the files?

- A. Azure Functions
- B. Microsoft SQL Server Integration Services (SSIS)
- C. PolyBase
- D. Azure Data Factory

Answer: C

NEW QUESTION 221

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questons and Answers in PDF Format

DP-900 Practice Exam Features:

- * DP-900 Questions and Answers Updated Frequently
- * DP-900 Practice Questions Verified by Expert Senior Certified Staff
- * DP-900 Most Realistic Questions that Guarantee you a Pass on Your FirstTry
- * DP-900 Practice Test Questions in Multiple Choice Formats and Updatesfor 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The DP-900 Practice Test Here](#)